Plaintiffs' Exhibit 22

Message

From:	Jim Giles [jimgiles@google.com]
Sent:	8/11/2017 4:52:34 PM
To:	Nitish Korula [nitish@google.com]; Rita Ren [rren@google.com]
CC:	Martin Pál [mpal@google.com]; Vahab Mirrokni [mirrokni@google.com]; Michael Hopkins [mhop@google.com]; Max Lin [whlin@google.com]; Max Loubser [maxl@google.com]; Balasubramanian Sivan [balusivan@google.com]; renatoppl [renatoppl@google.com]; Cyrille Berliat [cberliat@google.com]; Ali Radha [aradha@google.com]; Gang Wang [wanggang@google.com]
Subject:	Re: GOOGLE INTERNAL Copy of Q1 Programmatic OKR: margin manipulation investigation - Invitation to comment

For now, I like the fact that we still have last look for remnant -- that is one of the advantages of EB if publishers. #2 is the main thing I believe we should make sure is not a misperception, and I think using #1 to encourage publishers to push their partners into EBDA is worth doing.

On Fri, Aug 11, 2017 at 11:29 AM Nitish Korula <nitish@google.com> wrote:
I read their document as conflating two different issues:
1) AdX helps itself at the cost of AppNexus - this is the last look issue Martin brings up, and is indeed legitimate
2) AdX helps itself at the cost of lower revenue to publishers - this is because they don't know we do debt recollection from buyers / make buyers pay more than they otherwise would have. This is achieved in DRSv2 by non-truthful debt recollection, and in tDRS potentially by bundling with (parts) of RPO.

I *think* they're more concerned about the latter, though probably the former to some extent as well. I wonder if we explain to them the specific fact that buyer debt recollection results in higher publisher revenue through charging the buyer more, whether that would help?

(In particular, if my math is right, *even under their assumptions*, they make 10 pounds more under DRS than under the current world. But also, we tend to take higher margin on low-value queries, etc.)

Nitish
On Fri, Aug 11, 2017 at 11:12 AM, Rita Ren <rren@google.com> wrote:
If we make a small tweak on tDRS to not using the maximum remnant CPM in the revshare prediction, then we'll achieve what Martin mentioned at the end.

However, as long as our prediction is doing fairly well, under some scenarios we could in theory result in lower publisher overall revenue from DRS as many stated above. Bundling with RPO could help take us to a better position, but that would mean bundling up the opt-in for them together as well? I guess it would be a bit harder to explain this bundled feature to pubs, but a/b result could be very helpful.

On Fri, Aug 11, 2017 at 7:56 AM, Martin Pál <mpal@google.com> wrote:
My take is that the main culprit is not dynamic revshare, but rather the fact that AdX gets last look over AppNexus (assuming AppNexus demand is booked as price priority LIs in DFP).

AdX gets to pay high and win whenever AppNexus is present with a high CPM, and can pay low when AppNexus bids low. AppNexus in contrast can't reliably save money on queries where AdX bids low, because it doesn't know AdX bids.

This has fundamentally nothing to do with dynamic revshare -- dynamic revshare is just yet another way for AdX to exploit the last look advantage.

(Of course, dynamic revshare can be used without a last look advantage -- but doing so requires fairly accurate prediction of the opponent's bid).

On Fri, Aug 11, 2017 at 12:12 AM, Vahab Mirrokni <mirrokni@google.com> wrote:
We can improve messaging, but as Michael summarized, their concern comes from the fact some variants of DRS implementation can/may result in lower overall revenue. For example, if tDRS is implemented without any bundling with part of RPO - this could happen (in principle). Granted that since tDRS has additional ML for price prediction, we may still increase total revenue as well - but as we improve RPO price prediction further, this issue can show up again. This is something that we have had in mind, especially for tDRS. If we want to be share per-impression data for tDRS and prove this point, we may need to bundle with part of reserve pricing. If this is an important matter, when the time comes, we should do a similar "total revenue" analysis for tDRS (similar to what we did for DRS v2) and some what level of bundling with part of RPO is needed.

At the end of the day, we should be ready to show this on data (similar to total revenue analysis that we did for DRS v2 - showing that for majority of publishers we increased it). For DRS v2, I do think data would show that they don't lose in total revenue (if their declared min floors are not too off) - but I agree that communicating the details of pricing may not work. Question is how can we communicate that we can provide data that we are increasing the revenue without details of pricing scheme? We may agree with a "total revenue" metric with them & show them some data on that metric?

[btw Renato is OOO on vacation for some time, and I'll be OOO for the next 3 weeks.]

On Thu, Aug 10, 2017 at 11:44 PM, Jim Giles <jimgiles@google.com> wrote:
I feel like we should consider communicating with them that their assumptions are incorrect and that we are trying to act in their interest. Leaving the impression that we are not is harmful I think, regardless of what decision they make about turning it on or off. What do others think?

On Thu, Aug 10, 2017 at 9:23 PM Nitish Korula <nitish@google.com> wrote:
I think so.

Nitish

On Thu, Aug 10, 2017 at 9:13 PM, Michael Hopkins <mhop@google.com> wrote:
It sounds like the Guardian's core concern is that they want to maximize their overall revenue across indirect sources, and they're worried that DRS increases their AdX revenue but decreases their overall revenue.

I think their skepticism is reasonable -- you could imagine a version of DRS w/o buyer debt and that only took a lower rev share when a value cpm from a dfp remnant line item set the floor.

There are two properties of DRS that they didn't include in their analysis which can lead to increases in overall pub rev (although the effect of DRS on overall revenue can go either way):
1) More aggressive pricing against AdX buyers when they win below the floor
2) The AdX floor might not always represent a true opportunity cost (e.g. AdX open auction pricing rule).

The difficulty in emphasizing 1) is that it can open up questions about the mechanism, which we would like to avoid going into the details of,

We could perform an analysis where we look at the effect of DRS on the Guardian's overall revenue (treating DFP remenant line item value cpms as true opportunity costs), but it sounds like the Guardian would continue to be skeptical until seeing impression level data.

Once we launch tDRS would we feel comfortable including per query rev share in the bid level DT file?

Thanks,
Michael

On Thu, Aug 10, 2017 at 7:42 PM Max Lin <whlin@google.com> wrote:
We do not run the dynamic revshare scheme as the Guardian imagines. While we may not want to share the exact algorithm (given DRS is still evolving), can we emphasize that we are publisher's agent and we won't make product change that would decrease publisher's payout? Surfacing DRS as an opportunity and surface a/b result can provide data and back up our claim.

On Thu, Aug 10, 2017 at 6:06 PM Nitish Korula <nitish@google.com> wrote:
Thanks, Max. It's interesting, because they have multiple incorrect assumptions. In particular, we (a) Take a lower margin on higher-price queries and (b) Recollect debt from buyers. Did we try communicating this with them?

Thanks,
Nitish

On Thu, Aug 10, 2017 at 5:22 PM, Max Loubser (via Google Docs) <drive-shares-noreply@google.com> wrote:
Max Loubser has invited you to **comment on** the following document:

## GOOGLE INTERNAL Copy of Q1 Programmatic OKR: margin manipulation investigation

This is a document shared by The Guardian after their decision to turn off DRS.

Google Docs: Create and edit documents online.

Google Inc. 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA

You have received this email because someone shared a document with you from Google Docs.

--
Max Lin